

Distributed User Profiling via Spectral Methods

Dan-Cristian Tomozei
EPFL, Switzerland
dan-cristian.tomozei@epfl.ch

Laurent Massoulié
Technicolor, France
laurent.massoulie@technicolor.com

Abstract

User profiling is a useful primitive for constructing personalised services, such as content recommendation. In the present paper we investigate the feasibility of user profiling in a distributed setting, with no central authority and only local information exchanges between users. We compute a *profile vector* for each user (i.e., a low-dimensional vector that characterises her taste) via spectral transformation of observed user-produced ratings for items. Our two main contributions follow:

- i) We consider a low-rank probabilistic model of user taste. More specifically, we consider that users and items are partitioned in a constant number of classes, such that users and items within the same class are statistically identical. We prove that without prior knowledge of the compositions of the classes, based solely on few random observed ratings (namely $O(N \log N)$ such ratings for N users), we can predict user preference with high probability for unrated items by running a local vote among users with similar profile vectors. In addition, we provide empirical evaluations characterising the way in which spectral profiling performance depends on the dimension of the profile space. Such evaluations are performed on a data set of real user ratings provided by Netflix.
- ii) We develop distributed algorithms which provably achieve an embedding of users into a low-dimensional space, based on spectral transformation. These involve simple message passing among users, and provably converge to the desired embedding. Our method essentially relies on a novel combination of gossiping and the algorithm proposed by Oja and Karhunen.

Keywords Spectral Decomposition Random Matrix, Message Passing, Distributed Spectral Embedding, Distributed Recommendation System

1 Introduction

Recommendation systems have attracted much interest lately, mostly because of their relevance to core businesses of several major companies (e.g. Amazon, Netflix, Yahoo) who offer large catalogues of products to a vast user base. While the advertisement of highly popular items is straightforward, a significant portion of business stems from sales of only mildly popular items. The latter cannot be advertised indiscriminately, and must be recommended to the “right” users, through targeted recommendations. Such companies dispose of large storage and computational resources which enable a centralised computation of recommendations.

In this paper we take a different perspective on the problem of recommendation. Namely, we aim to develop strategies suited to *distributed* operation, where the burden of recommendation is not offloaded to the server, but is rather shared among the users. More specifically, we propose the following two-stage approach for generating recommendations:

- In the first stage, distributed algorithms assign coordinates (or *profiles*) to the users within a certain *profile space*, such that proximity in this space translates to proximity of user taste for content. We say that such algorithms perform *user profiling*.
- In the second stage, recommendations are obtained via simple and distributed algorithms which rely on the primitive of user profiling. We thereby avoid the need for complex machine learning techniques.

The performance of such an approach will depend heavily on the properties of the considered embedding of users in the profile space. For this reason, the focus in this paper is on the first stage of the process, i.e., *user profiling*. Namely, we argue that *spectral* profiling techniques retrieve hidden structure.

The techniques employed in a centralised setting for generating content recommendation are widely known under the generic name of “collaborative filtering”. They are typically implemented by a provider who wishes to offer a recommendation service to a large customer base. In such a setting, the information requested from the customers (or users) is typically related both to their identity (via the registration procedure) and to their taste (via the opinions they express regarding the items).

It is not clear to which extent identity information characterises user taste. Moreover, the nature of such information gives rise to privacy concerns. On the contrary, the opinions that users express about items constitute the truly relevant data for solving the problem. For this reason, we advocate a purely agnostic approach to recommending content, which does not use information about the real identity of users, or the nature of content.

Opinions are expressed in the form of ratings assigned by a user to the items she has already purchased. Ratings characterise the satisfaction of a user with respect to a specific item. They are discrete and range from a lowest to a highest value (e.g., number of stars). In particular, the mere fact that a user has consumed or not a specific item can be regarded as a *binary rating*. In this paper we consider the latter form of rating.

Since the number of items on offer from the provider is overwhelmingly large, the vast majority of users only consume a small fraction of items. Hence, for a typical user, only a small number of ratings are known. In the case of binary ratings, if an item has not been consumed it does not necessarily follow that the user dislikes it. It is possible that the user is simply unaware of the item’s existence. Hence, in this case we cannot distinguish between missing ratings and disliked items.

A recent illustration of the possible machine learning techniques and of the corresponding performance comes from the Netflix prize competition [12]. The goal of the competition was to design an algorithm that, when trained on a data set made publicly available by Netflix, would manage to improve prediction accuracy by 10% (measured via Root Mean Squared Error) compared to the proprietary Cine-match algorithm. The designers of such an algorithm [9] were awarded a prize of \$1M three years after

the start of the competition. It is important to note that the last two years of the three were spent trying to improve the gain in prediction from 8.42% to 10.04%.

The strenuous advancement of the Netflix prize suggests the existence of an important obstacle in the way of achieving high prediction accuracy. A possible explanation is given in the study conducted by Amatriain et al. [1], which showed that when presented with a movie title several times, users provide inconsistent ratings. Hence, there exists an implicit noise in any collection of ratings due to the fact that human taste is variable and not easily quantified. The authors [1] used the RMSE to characterise the distance between two sets of ratings assigned by the same users to the same movies. They found RMSE values comparable to the winning entry of the Netflix prize. These results suggest that user preference has a significant random component. As such, it is reasonable to consider a probabilistic model of user taste. Throughout this paper we make the following

Assumption 1. *The taste of each user is characterised by a certain probability distribution defined on the set of all possible ratings for the set of items. A user's observed ratings are obtained by sampling her corresponding multi-dimensional distribution.*

Denote the set of users by \mathcal{U} and the set of items by \mathcal{F} . A natural representation for the observed binary rating information is as follows. Consider a rectangular matrix $S \in \{0, 1\}^{|\mathcal{U}| \times |\mathcal{F}|}$ which we call *the rating matrix*. Each row corresponds to a user and each column corresponds to an item. An entry S_{ui} corresponds to the user-item pair $(u, i) \in \mathcal{U} \times \mathcal{F}$. The entry holds 1 if user u has purchased item i , and 0 otherwise. As previously stated, most of the zero entries in the rating matrix correspond to cases in which a user has not considered purchasing a specific item. According to Assumption 1, each row of matrix S is partially observed realisation of a $\{0, 1\}$ random $|\mathcal{F}|$ -dimensional vector. Relying only on matrix S , we need to assign profiles to the users, such that users with similar taste have similar profiles.

The instances of the problem we consider are extremely large, it is not uncommon to have a user base \mathcal{U} of the order of millions and a catalogue of items \mathcal{F} of the order of tens of thousands. In the case of binary ratings, simply representing the probability distribution in Assumption 1 for a single user requires an exponential amount of memory $2^{|\mathcal{F}|}$. We are thus constrained to consider very simple approximations of such probability laws, for the sake of computational tractability.

Like most proposed models of user taste found in the literature (e.g., [7]), we consider a low-rank model of user taste. More specifically, we make the following

Assumption 2. *Each entry S_{ui} of the rating matrix is given by a Bernoulli random variable of parameter \bar{S}_{ui} . The matrix $\bar{S} = (\bar{S}_{ui})$ has rank K , where $K \ll |\mathcal{U}|$ and $K \ll |\mathcal{F}|$.*

Hence, we consider a low-rank *probabilistic* model of user taste, as opposed to a deterministic one [7].

In Section 3 we propose a user profiling technique based on the Singular Value Decomposition of matrix S . For a probabilistic model of user taste satisfying Assumption 2, and under further weak statistical assumptions, we prove that a simple voting scheme among users with similar profiles manages to produce accurate recommendations for most of the items with high probability. Furthermore, we use actual movie ratings to compute the profiles of anonymous users of the Netflix system. We observe that users with similar profiles have similar taste in movies.

Motivated by the ability to recover hidden structure of the spectral techniques, in the second part of this paper (Section 4), we design a distributed algorithm that computes individual spectral profiles based on local exchanges among users. We prove almost sure convergence of the algorithm and provide evaluations on a synthetic trace. We conclude in Section 5.

2 Related Work

Keshavan et al. [7] consider the problem of low rank matrix completion. They show that for a constant rank $r = O(1)$ “well-behaved” matrix, it is sufficient to have $\Omega(N \log N)$ revealed entries in order to be

able to achieve exact matrix reconstruction. For a square matrix, this corresponds to an average degree of $\Omega(\log N)$, just as in our result. The drawback to a direct application of this result to the user taste prediction problem is the fact that the sought matrix is deterministic. Implicitly, it would mean that user ratings are deterministic, and that rating matrices are low rank. In contrast, the models we consider are probabilistic.

In Section 3 we consider a low-rank probabilistic model of user taste. Users are partitioned into classes, such that users within the same class are statistically identical. We show (Theorems 1 and 2) that the profile vectors corresponding to each user computed via spectral methods are clustered around distinct points corresponding to the classes. In this respect, our results are related to spectral clustering. There is a vast literature on the topic of spectral clustering, of which we now give a brief overview.

Among the most relevant work, Ng et al. [13] propose a clustering algorithm based on spectral decomposition. Our results provide a comprehensive analysis by giving conditions under which the underlying partition into classes is retrieved exactly.

In [4], Dasgupta et al. propose an algorithm based on iterative splitting of groups into two subgroups. In contrast, we obtain the desired groups in one go. In [10], McSherry proposes a different clustering method based on projections onto the column space of the original matrix.

In both [4] and [10], the probabilistic model of the original matrix is akin to ours (which extends the classical “planted partition” model). In contrast, we establish our results under far less stringent conditions on the average degree of the original matrix. Namely, we require an average degree of order $\Omega(\log(N))$ while they require an order of $\Omega(\log(N)^6)$.

The recent paper [15] by Shi et al. discusses rationales for choosing which eigenvectors to use when performing spectral clustering. This issue is to a large extent complementary to the ones we address in this paper. We could rely on [15] to specify which eigenvectors to keep in our profiling context.

In Section 4.1 we propose a method for computing the eigenvectors of the adjacency matrix of a graph in a distributed manner. A variant of the method was briefly described in [17]. Eigenvector extraction is the object of Oja’s algorithm [14]. This basic algorithm was refined by Borkar and Meyn [2]. None of these approaches is distributed however. Our contribution in Section 4.1 consists precisely in augmenting these methods to make them distributed.

A significant contribution towards computing the top k eigenvectors of a symmetric weighted adjacency matrix in a distributed fashion was brought by Kempe and McSherry in [6]. The setting is similar to the one we consider in Section 4.1. The authors give bounds on the required running time of their algorithm. Due to the fact that we explicitly introduce noise in our iterations, obtaining such bounds is more difficult in our case. However, while the algorithm in [6] performs an explicit orthonormalisation within the main loop, which requires a global synchronization point, our approach facilitates an asynchronous implementation. The normality of the vectors computed using our algorithm is a direct consequence of Oja’s algorithm. Furthermore, the algorithm we analyze in Section 4.1 requires simpler computations at each of the nodes compared to [6].

A similar approach to the one presented in this paper has been taken in a recent publication [8]. The authors aim to determine the eigenvectors of a deterministic matrix based on random sparse observations. They derive useful bounds on convergence time. However, the gossiping stage in their proposed algorithm is treated as a “black box”. We explicitly construct an algorithm that incorporates two stages: gossiping, performed on a faster time scale, and Oja’s method, performed on a slower time scale. Moreover, we explicitly determine multiple eigenvectors, whereas the authors of [8] focus on determining a single eigenvector, and argue that the extension can be achieved. Finally, we propose an asynchronous algorithm. We show on synthetic data that the latter determines the desired eigenvectors.

3 Spectral Recovery of Probabilistic Taste

We begin by analysing a simple setting in which the observations consist of measures of similarity between users. We prove that a profiling technique based on the spectral decomposition of the square matrix regrouping the observed similarities between pairs of users successfully recovers hidden structure.

We apply these findings to the case in which ratings of items by users are observed. We show that a simple distributed voting algorithm provides asymptotically accurate predictions for most items.

Finally, we observe the benefits of spectral profiling on a real trace.

We make use of notations described in Table 1. Unless otherwise indicated, all vectors are column vectors.

A'	The transposition of matrix A
x'	The transposition of column vector x
e	The all-ones column vector
$\text{diag}(\alpha)$	The $K \times K$ diagonal matrix having the elements on the main diagonal given by the K -dimensional vector α
$\ y\ _\alpha$	The α -norm of vector y , where $0 < \alpha < 1$ and y are column vectors with the same dimension, $\ y\ _\alpha := \sqrt{\sum_k \alpha_k y_k^2}$
$y(k), y_k$	The k -th element of column vector y

Table 1: General notations

3.1 Similarity-based Profiling

Denote by $N = |\mathcal{U}|$ the number of users. Let us consider in this first stage that we are given partial observations of user taste similarity in the form of a symmetric matrix

$$A \in \{0, 1\}^{N \times N}.$$

Namely, for any two users u, v the elements $A_{uv} = A_{vu}$ take value 1 if users u and v have been evaluated as similar, and value 0 if the users are deemed dissimilar, or if the similarity between the two users has not been evaluated. By convention $A_{uu} = 0$.

We propose the following *spectral* representation of users based on these partial similarity observations. For some fixed dimension L , extract the *normalised* eigenvectors x_1, \dots, x_L corresponding to the L largest magnitude eigenvalues of matrix A . We define the *profile space* as \mathbb{R}^L . In the profile space, to each user u there corresponds a scaled row vector $\sqrt{N}z'_u$, where

$$z'_u = (x_1(u), \dots, x_L(u)).$$

We refer to this vector as the *profile of user u* . The scaling factor \sqrt{N} is introduced to compensate for the fact that the eigenvectors of the $N \times N$ matrix A are taken of norm 1. In what follows, we propose a simple probabilistic user taste model for which this spectral representation of users enables us to retrieve hidden structure.

3.1.1 Statistical Model

We assume the following probabilistic model of user taste: The N users are partitioned into K classes C_1, \dots, C_K , such that users within the same class are statistically identical. We denote the size of class C_k , $1 \leq k \leq K$, by $|C_k| = \alpha_k N$, for fixed $\alpha_k > 0$, which are such that $\sum_k \alpha_k = 1$. For any user $u \in \mathcal{U}$ we denote by $k(u)$ her unique corresponding class.

u, v	Indices referring to users
k, ℓ	User profile indices
$k(u)$	Index of the profile containing user $u : u \in C_{k(u)}$
$\{C_k\}_{k=1}^K$	Partition of users into K unknown disjoint profiles
α	Vector grouping fractions of users per profile: $ C_k = \alpha_k N$
B	Constant unknown $K \times K$ matrix of probabilities
$p = \frac{\omega}{N}$	Probing probability
L	Dimension of profile space

Table 2: Notations and conventions

Each pair of classes $1 \leq k, \ell \leq K$ is characterised by probabilities $b_{k\ell} = b_{\ell k}$ as follows. User $u \in C_k$ and user $v \in C_\ell$ are similar with probability $b_{k\ell}$ and dissimilar with probability $1 - b_{k\ell}$. Moreover, for all pairs of users, their similarity is observed with probability $p = \frac{\omega}{N}$, where ω is a parameter of the model. The similarity of unobserved pairs of users is set to 0 by default.

Equivalently, for all ordered pairs of users $u < v$, the observed similarities $A_{uv} = A_{vu}$ are the outcome of independent Bernoulli random variables of parameters $pb_{k(u)k(v)}$. As previously stated, the diagonal elements are all null, $A_{uu} = 0$.

Given the observed similarity matrix A , without knowledge of the $K \times K$ profile similarity matrix $B = (b_{k\ell})_{k,\ell}$, we wish to recover the partition of users in the unknown classes C_k using their spectral representation. Below we provide sufficient scaling assumptions for which such recovery is possible.

In Table 2 we summarise the notations we have introduced so far.

3.1.2 Scaling Assumptions

Let us describe the dependence of the various parameters of the model on the number of users.

We have made the Assumption 2, namely that our model is low-rank. Thus, we consider that the number of classes K , as well as the fraction α_k of users in each class, and the similarity probability matrix $B = (b_{k\ell})_{k,\ell}$ are constant as the number of users in the system grows.

However, we assume that the probing probability p vanishes as the number of users grows. Namely, ω goes to infinity slower than N ,

$$\omega \rightarrow_N \infty, \omega = o(N). \quad (1)$$

Hence, each user probes on average a fraction $p = \frac{\omega}{N} \rightarrow 0$ of users with whom she evaluates taste similarity. Equivalently, matrix A can be regarded as the adjacency matrix of a random graph on the set of users having average degree of $\Theta(\omega)$.

3.1.3 Hidden Structure Recovery

Theorem 1 below states that under mild conditions, for large N , the spectral representations $\sqrt{N}z'_u$ of users are clustered according to their respective classes C_k . Consider the vector

$$\alpha := (\alpha_k)_{k=1}^K.$$

The α -norm of a K -dimensional vector t is $\|t\|_\alpha^2 = \sum_k \alpha_k t_k^2$ (see Table 1). Define the following constant matrix

$$M := (b_{k\ell}\alpha_\ell)_{1 \leq k, \ell \leq K}.$$

Before stating the theorem, we introduce the following conditions:

The dimension of the profile space is upper bounded by $L \leq \text{rank}(M)$. (2a)

The L largest magnitude eigenvalues of M have distinct absolute values. (2b)

The corresponding eigenvectors normalised under the α -norm, y_1, \dots, y_L , satisfy:

$$t'_k \neq t'_\ell, 1 \leq k \neq \ell \leq K,$$

where $t'_k := (y_1(k), \dots, y_L(k))$. (2c)

$\omega \geq C \log(N)$, for some absolute constant C . (2d)

Theorem 1. *Under assumptions (2), with probability $1 - o(1)$, a fraction of $1 - o(1)$ users u is such that $\|\sqrt{N}z'_u - t'_{k(u)}\| = o(1)$, where $k(u)$ denotes the class of u . That is to say, most users have their (scaled) profile vector close to a fixed vector characteristic of their class.*

Each vector t'_k corresponds to a class C_k . If two classes k, ℓ were such that $t_k = t_\ell = t$, the theorem ensures that the profiles of users of both classes would be grouped around the constant vector t in the profile space. Hence, it would be impossible to distinguish between the users of the two classes based solely on their profile vectors. Condition (2c) ensures that

$$\|t'_k - t'_\ell\| = \Omega(1),$$

thus guaranteeing the ability to distinguish between distinct classes for a large enough number of users.

We prove that eigenvectors corresponding to non-zero eigenvalues of matrix A can be used to recover the hidden classes. Condition (2a) ensures selection of a suitable dimension L . We impose the technical condition (2b) for presentation ease.

Condition (2d) gives a lower bound of $\log N$ on order of the required average user neighbourhood size ω . The theorem states that for ω growing at least as fast as $\log N$, the initial partition in profiles C_k can be recovered with high probability for almost all users.

Remark 1. *Concerning condition (2d), it is plausible that an even lower requirement for ω (i.e. constant) suffices. Such a case has been explored in the context of bounding the second eigenvalue of the adjacency matrix of a sparse random graph to $O(\sqrt{\omega})$, by removing high degree outlier nodes from the graph (see [5]).*

Remark 2. *Other flavours of matrices than the adjacency matrix A could be considered for spectral analysis (e.g. Laplacian matrix, normalised Laplacian matrix). We do not address either of these scenarios in the present work.*

We now give the main steps in the proof of Theorem 1. The auxiliary lemmas are proved in the appendix. Consider the matrix

$$\bar{A} := (pb_{k(u)k(v)})_{u,v},$$

which, according to our model, is equal to the expectation $\mathbb{E}A$ of the partially observed similarity matrix A . We can write $A = \bar{A} + Q$, where $Q := A - \bar{A}$.

The theorem relies on the fact that the block matrix \bar{A} imposes the eigenvalues and eigenvectors of A , while the perturbation matrix Q has little influence therein, as follows from the lemmas below:

Lemma 1. *The top L largest magnitude eigenvalues of \bar{A} have distinct absolute values and are order of $\Theta(\omega)$. The normalised eigenvectors $(\bar{x}_\ell)_{\ell=1}^L$ corresponding to these eigenvalues are constant on indices corresponding to each user class. Specifically, using the y_ℓ defined in (2c), we can write*

$$\bar{x}_\ell(u) = \frac{y_\ell(k(u))}{\sqrt{N}}, \forall u \in \mathcal{U}, 1 \leq \ell \leq L. \quad (3)$$

Use the following ordering of the eigenvalues of \bar{A} and A :

$$\begin{aligned} |\bar{\lambda}_1| &> |\bar{\lambda}_2| > \dots > |\bar{\lambda}_L| \\ |\lambda_1| &\geq |\lambda_2| \geq \dots \geq |\lambda_L|. \end{aligned} \quad (4)$$

We also denote by \bar{x}_k and x_k the corresponding *normalised* eigenvectors.

To control the influence of the perturbation matrix Q we use the following

Lemma 2. *Consider a square $N \times N$ symmetric 0-diagonal random matrix A such that its elements $A_{ij} = A_{ji}$ are independent Bernoulli random variables with parameters $\mathbb{E}A_{ij} = p_{ij}\omega N^{-1}$, where the p_{ij} are constant and $\omega = \Omega(\log N)$. Then with high probability the spectral radius of the matrix $A - \mathbb{E}A$ satisfies the upper bound $\rho(A - \mathbb{E}A) \leq O(\sqrt{\omega})$.*

Denote by $D = \text{diag}(b_{k(u)k(u)} \frac{\omega}{N}, 1 \leq u \leq N)$. By application of Lemma 2 to matrix A , we get that the spectral radius of $Q_0 := A - \mathbb{E}A$ is upper bounded by $O(\sqrt{\omega})$ with high probability. Since $\rho(D) \leq O(\frac{\omega}{N})$, we have that the spectral radius of $Q = Q_0 - D$ is also upper bounded by $O(\sqrt{\omega})$ with high probability.

The previous two lemmas are instrumental in proving that \bar{A} and A have the same spectral structure:

Lemma 3. *Using the ordering (4) for the eigenvalues of \bar{A} and A , it holds that for all $1 \leq k \leq K$*

$$||\lambda_k| - |\bar{\lambda}_k|| \leq O(\sqrt{\omega}) \quad \text{whp}, \quad (5)$$

$$\sin(\widehat{x_k, \bar{x}_k}) \leq O(\omega^{-1/4}) \quad \text{whp}. \quad (6)$$

We conclude by an application of Tchebitchev's inequality. Lemma 3 shows that with high probability we have $\|x_\ell - \bar{x}_\ell\|^2 \leq O(\omega^{-1/2})$, for all $1 \leq \ell \leq L$. Condition (2c) guarantees that $\|t'_k - t'_\ell\| = \Omega(1)$.

By Tchebitchev's inequality, for any constant $a > 0$, and for a user U chosen uniformly at random, denoting by $k(U)$ its class, it holds that:

$$\mathbb{P}(\|\sqrt{N}z'_U - t'_{k(U)}\| \geq a) \leq \frac{1}{N} \sum_{u=1}^N \frac{\|\sqrt{N}z'_u - t'_{k(u)}\|^2}{a^2} \stackrel{(3)}{=} \sum_{\ell=1}^L \frac{\|x_\ell - \bar{x}_\ell\|^2}{a^2} = O(a^{-2}\omega^{-1/2}).$$

Thus we will be able to conclude the result of the theorem if we can find an a such that $a = o(1)$ and $a^{-2}\omega^{-1/2} = o(1)$. It is easy to see that for instance $a = \omega^{-1/6}$ satisfies these conditions. \square

3.2 Application: Extension to Content Recommendation

Let us now consider a seemingly different scenario. Denote again by $N = |\mathcal{U}|$ the number of users and by $F = |\mathcal{F}|$ the number of items. Assume without loss of generality that $N > F$. We can write $F = \gamma N$, with $0 < \gamma < 1$. We consider the rectangular observed rating matrix

$$S \in \{0, 1\}^{N \times F},$$

where $S_{ui} = 1$ if user u has rated and liked item i and 0 otherwise. If an entry S_{ui} is null, it is not necessarily true that user u dislikes item i (the item might have simply not been rated).

We propose the following representation of users based on this collected information. For some dimension L , extract the L *normalised left singular vectors* x_1, \dots, x_L of S , corresponding to its L largest singular values. Like in the previous subsection, consider the L -dimensional profile space \mathbb{R}^L , in which we associate a scaled row vector $\sqrt{N}z'_u$ to user u , where

$$z'_u = (x_1(u), \dots, x_L(u)).$$

The scaling by \sqrt{N} is again due to the fact that the singular vectors are normalised.

Recall that the Singular Value Decomposition (SVD) of a rectangular real matrix $S = X\Sigma Y$ is always well defined, where matrices $X \in \mathbb{R}^{N \times F}$ and $Y \in \mathbb{R}^{F \times F}$ are unitary and matrix $\Sigma \in \mathbb{R}^{F \times F}$ is positive diagonal.

Again, we claim that users with similar taste will be mapped to close-by locations in the profile space. More specifically, we show that local voting in the profile space provides users with accurate recommendations for most of the items. In what follows, we give a probabilistic model of user taste and scaling assumptions for which we prove these claims.

3.2.1 Statistical Model

We consider a probabilistic model of user taste similar to the one in Section 3.1.1. The N users are partitioned into K disjoint classes C_1, \dots, C_K and the F items are partitioned into K' disjoint classes $D_1, \dots, D_{K'}$. Users and items in the same class are statistically identical. The size of user class C_k is denoted by

$$|C_k| = \alpha_k N, \quad 1 \leq k \leq K,$$

while the size of item class $D_{k'}$ is denoted by

$$|D_{k'}| = \beta_{k'} F, \quad 1 \leq k' \leq K'.$$

The $(\alpha_k)_k$ and $(\beta_{k'})_{k'}$ are strictly positive and sum to 1:

$$\sum_k \alpha_k = \sum_{k'} \beta_{k'} = 1, \quad \alpha_k > 0, \beta_{k'} > 0.$$

For any user $u \in \mathcal{U}$ we denote by $k(u)$ her unique corresponding user class, and for any item $i \in \mathcal{F}$ we denote by $k'(i)$ its unique corresponding item class.

Pairs of user and item classes $(C_k, D_{k'})$, $1 \leq k \leq K, 1 \leq k' \leq K'$, are characterised by probabilities $r_{kk'}$ in the following way: Any user $u \in C_k$ likes any item $i \in D_{k'}$ with probability $r_{kk'}$ and dislikes it with probability $1 - r_{kk'}$. For all user-item pairs $(u \in \mathcal{U}, i \in \mathcal{F})$, u decides to rate i with probability $p = \omega N$. Equivalently, any element S_{ui} of the observed rating matrix S is obtained by drawing a Bernoulli random variable of parameter $p \cdot r_{k(u)k'(i)}$.

Given the observed rating matrix S , without knowledge of the $K \times K'$ affinity matrix $R := (r_{kk'})_{k,k'}$, we wish to recover the partition of users in the K classes C_k by making use of their spectral representation. We provide sufficient scaling conditions in what follows.

We summarise the notations we use in Table 3.

3.2.2 Scaling Assumptions

We make again Assumption 2 and take the number of classes of users K and of items K' to be constant with N . We assume that the fraction of users in each user class α_k , the fraction of items in each item class $\beta_{k'}$, the ratio between the number of items and the number of users $\gamma = \frac{F}{N}$, as well as the class affinity matrix $R = (r_{kk'})_{k,k'}$ are constant with respect to N . Thus, the class sizes, as well as the total number of items grow linearly with N .

We assume that the rating probability p vanishes as the number of users N grows to infinity. Specifically, we again impose condition (1) on parameter ω . Note that the expected number of items rated by a user is order of $\Theta(\omega)$.

Let us now formulate and prove an extension of Theorem 1 that we can apply in this setting.

N	Number of users
$F = \gamma N$	Number of items
u, v	Indices referring to users
i, j	Indices referring to items
k, ℓ	User profile indices
k', ℓ'	Item class indices
$k'(i)$	Index of the class containing item $i : i \in D_{k'(i)}$
$\{C_k\}_{k=1}^K$	Partition of users into K disjoint classes
$\{D_{k'}\}_{k'=1}^{K'}$	Partition of items into K' disjoint classes
α	Vector grouping fractions of users per class: $ C_k = \alpha_k N$
β	Vector grouping fractions of items per class: $ D_{k'} = \beta_{k'} F$
R	Constant unknown $K \times K'$ affinity matrix

Table 3: Preferred indices

3.2.3 Content Recommendation

Consider vectors

$$\alpha = (\alpha_k)_{k=1}^K, \beta = (\beta_{k'})_{k'=1}^{K'},$$

and define the following constant square matrix

$$G := R \text{diag}(\beta) R' \text{diag}(\alpha) \in \mathbb{R}^{K \times K}.$$

We impose the following conditions, similar to (2):

$$\text{The dimension of the profile space is upper bounded by } L \leq \text{rank}(G). \quad (7a)$$

$$\text{The } L \text{ largest eigenvalues of matrix } G \text{ are distinct.} \quad (7b)$$

The corresponding eigenvectors $\{g_\ell\}_{\ell=1}^L$ normalised for the α -norm satisfy:

$$\chi_k \neq \chi_\ell, 1 \leq k < \ell \leq L,$$

$$\text{where } \chi_k = (g_1(k), \dots, g_L(k)). \quad (7c)$$

$$\omega \geq C \log N \text{ for some constant } C. \quad (7d)$$

We can now state the following

Theorem 2. *Under conditions (7), with probability $1 - o(1)$, a fraction of $1 - o(1)$ users u is such that $\|\sqrt{N}z'_u - \chi_k\| = o(1)$. That is to say, most users have their scaled profile vector close to a fixed vector corresponding to their class.*

Before giving the main points of the proof of Theorem 2, let us explain how this seemingly distinct setup can be mapped to the previous one. Define the following transformation which produces a square matrix:

$$\tau : S \mapsto A = \begin{bmatrix} 0 & S \\ S' & 0 \end{bmatrix} \in \mathbb{R}^{(N+F) \times (N+F)},$$

where S' denotes the transposition of matrix S .

The spectrum of such a matrix is symmetrical (i.e. if σ is an eigenvalue of A , then so is $-\sigma$). Furthermore, the absolute values of the eigenvalues of A are the singular values of S and its singular components

are determined by the eigenvectors of A . To see this, consider an eigenvalue σ of A and its corresponding eigenvector $\zeta = \begin{bmatrix} x \\ y \end{bmatrix}$, with $x \in \mathbb{R}^{N \times 1}$ and $y \in \mathbb{R}^{F \times 1}$. Since $A\zeta = \sigma\zeta$, we can write:

$$A \begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} Sy \\ S'x \end{bmatrix} = \sigma \begin{bmatrix} x \\ y \end{bmatrix} \text{ and } A \begin{bmatrix} x \\ -y \end{bmatrix} = \begin{bmatrix} -Sy \\ S'x \end{bmatrix} = -\sigma \begin{bmatrix} x \\ -y \end{bmatrix},$$

and thus $x'Sy = y'S'x = \sigma\|x\|^2 = \sigma\|y\|^2$. Then, we have that x is a left singular vector, and that y is a right singular vector for matrix S . For more details see for instance [16].

We can now prove Theorem 2. We make the following notation:

$$\bar{S} := (pr_{k(u)k'(i)})_{u,i} = \mathbb{E}S.$$

The following Lemmas characterise the structure of the singular decomposition of \bar{S} and S . They show that the two matrices have the same spectral structure. The proofs are found in the appendix.

Lemma 4. *For $L \leq K'$, the top L largest singular values of \bar{S} are distinct and of order $\Theta(\omega)$. The normalised left-singular vectors $(\bar{x}_\ell)_{\ell=1}^L$ corresponding to these singular values are constant on indices corresponding to each user class. Specifically, using the g_ℓ defined in (7c), we can write*

$$\bar{x}_\ell(u) = \frac{g_\ell(k(u))}{\sqrt{N}}, \quad \forall u \in \mathcal{U}, \quad 1 \leq \ell \leq L. \quad (8)$$

We denote the singular values of \bar{S} and S by

$$\begin{aligned} \bar{\sigma}_1 &> \bar{\sigma}_2 > \dots > \bar{\sigma}_L > 0 \\ \sigma_1 &\geq \sigma_2 \geq \dots \geq \sigma_L > 0. \end{aligned} \quad (9)$$

and the corresponding left and right normalised singular vectors by $\bar{x}_k, \bar{y}_{k'}, x_k$, and $y_{k'}$.

Lemma 5. *For all $1 \leq k \leq K$*

$$|\sigma_k - \bar{\sigma}_k| \leq O(\sqrt{\omega}) \quad \text{whp}, \quad (10)$$

$$\sin(\widehat{x_k, \bar{x}_k}) \leq O(\omega^{-1/4}) \quad \text{whp}. \quad (11)$$

To conclude, let us use a similar argument to the one in the proof of Theorem 1. Lemma 5 shows that with high probability we have $\|x_\ell - \bar{x}_\ell\|^2 \leq O(\omega^{-1/2})$, for all $1 \leq \ell \leq L$. Condition (7c) guarantees that $\|\chi'_k - \chi'_\ell\| = \Omega(1)$.

Again, by Tchebitchev's inequality, for any constant $a > 0$, and for a user U chosen uniformly at random, it holds that:

$$\mathbb{P}(\|\sqrt{N}z'_U - \chi'_{k(U)}\| \geq a) \leq O(a^{-2}\omega^{-1/2}). \quad (12)$$

We conclude the proof by choosing again $a = \omega^{-1/6}$. \square

3.2.4 Characterising Performance of a Simple Voting Algorithm

Let us now analyse a simple recommendation algorithm that relies on local voting in the profile space. We have defined for each user u a scaled L -dimensional profile vector which we denoted by $\sqrt{N}z'_u$. In Section 4 we propose a method for computing such profile vectors in a distributed fashion based solely on local information exchange.

Say we wish to characterise the taste of user u for item i . Consider a fixed constant $d > 0$. We define the d -vicinity of u as the set of users that have profile vectors at a Euclidean distance of at most d from $\sqrt{N}z'_u$ in the profile space. More formally,

$$V_d(u) := \{v \in \mathcal{U} : \|z_u - z_v\|_2 \leq \frac{d}{\sqrt{N}}\}.$$

A vote is run among users in $V_d(u)$ to collect their appreciation for item i . The estimated preference $\hat{p}_{u,i}$ of user u for item i is defined as the proportion of users in $V_d(u)$ that have given a rating of 1 to i , that is

$$\hat{p}_{u,i} = \frac{1}{|V_d(u)|} \sum_{v \in V_d(u)} S_{vi}. \quad (13)$$

The following proposition guarantees that, for well chosen d and as the number of users grows to infinity, the voting algorithm above produces accurate recommendations for most users and most items.

Proposition 1. Define $\delta := \min_{1 \leq k < \ell \leq L} \|\chi_k - \chi_\ell\|_2$, where χ_k are defined in (7c). If $d < \delta/2$, then with high probability, for a fraction of $(1 - o(1))$ of users u and a fraction $(1 - o(1))$ of items i , the estimator $\hat{p}_{u,i}$ defined in (13) obtained via the voting procedure is of the form

$$\hat{p}_{u,i} = \Theta\left(\frac{\omega}{N}\right) (r_{kk'} + o(1)).$$

Proof. For any user u denote the distance from her scaled profile vector to the vector $\chi_{k(u)}$ corresponding to her class by $d_0(u) := \|\sqrt{N}z'_u - \chi_{k(u)}\|_2$. For some positive D , we call D -good users those users u having $d_0(u) < D$.

A simple application of (12) for constant D gives

$$\mathbb{P}(d_0(U) \geq D) \leq O(\omega^{-1/2}),$$

where U is a user taken uniformly at random. Denote by $Y(D)$ the number users that have their scaled profile vector at distance larger than D from the vector χ_k corresponding to their class. Then it follows that,

$$\mathbb{E}Y(D) \leq O\left(\frac{N}{\sqrt{\omega}}\right) = o(N).$$

Thus there exists a fraction $(1 - o(1))$ of D -good users.

Let ϵ be a fixed constant such that $0 < \epsilon < d$. We consider $(d - \epsilon)$ -good users u of class k , that is with their scaled profile vectors at distance at most $d - \epsilon$ of χ_k (i.e., $d_0(u) < d - \epsilon$).

User $u \in C_k$ wishes to obtain a recommendation for a certain item i and runs a vote among the users in $V_d(u)$ (i.e., within the ball of radius d centred at her profile vector) as described above. This ball contains the set of ϵ -good users of class k , which we denote by $V_\epsilon^k := \{v \in C_k : d_0(v) < \epsilon\}$. Since ϵ is a constant, the number of users outside of $\cup_\ell V_\epsilon^\ell$ (i.e., that are not ϵ -good) is at most $o(N)$. We can write

$$|V_\epsilon^k| = |C_k| - |\bar{V}_\epsilon^k| = \alpha_k N + o(N),$$

where \bar{V}_ϵ^k denotes the set of users of class k that are not ϵ -good (i.e., $\bar{V}_\epsilon^k = C_k \setminus V_\epsilon^k$).

For any $(d - \epsilon)$ -good user u of class k , by construction we have that $V_\epsilon^\ell \cap V_d(u) = \emptyset$, for all $\ell \neq k$. Thus, the set of users in $V_d(u)$ which are not of class k is upper-bounded by

$$|V_d(u) \setminus C_k| \leq |\cup_{\ell \neq k} \bar{V}_\epsilon^\ell| \leq o(N).$$

Moreover, since $V_\epsilon^k \subset V_d(u)$,

$$|V_d(u) \cap C_k| \geq |V_\epsilon^k| = \alpha_k N + o(N).$$

Thus, most of the users of class C_k are present in $V_d(u)$, and at most $O(N/\sqrt{\omega})$ outliers pollute the voting scheme. In what follows we show that the influence of the outliers vanishes asymptotically.

Denote by W the set of items having received more than f votes from the $O(\frac{N}{\sqrt{\omega}})$ outliers in $V_d(u)$. Then we have that with high probability

$$|W| \leq \frac{\sqrt{\omega}N}{f}.$$

Take $f = o(\omega)$, but such that $|W| = o(N)$. For example, pick $f = \omega^{3/4}$.

For any item $i \notin W$ (there is a fraction $(1 - o(1))$ of such items), denote its class by $k' = k'(i)$. Then we can write

$$\hat{p}_{u,i} = \frac{1}{\alpha_k N + o(N)} \left(\sum_{v \in V_d(u) \cap C_k} X_v + O(f) \right),$$

where the X_v are independent Bernoulli random variables identically distributed with parameter $pr_{kk'} = \frac{\omega}{N} r_{kk'}$. They correspond to the votes of users in the same class C_k as u that belong to $V_d(u)$. Since the number of such users is at least $\alpha_k N - O(\frac{N}{\sqrt{\omega}}) = \Theta(\alpha_k N)$, by application of a Chernoff bound we obtain that with high probability

$$\hat{p}_{u,i} = \Theta\left(\frac{\omega}{N}\right) \left(r_{kk'} + \underbrace{O(f\omega^{-1})}_{o(1)} \right),$$


which concludes the proof. \square

3.3 Spectral Profiling in Practice


In this section we evaluate the benefits of spectral techniques on a real trace provided by Netflix. The Netflix data set contains about 10^8 user ratings for 17,770 movies by 480,000 users. The ratings are given in the form of an integer number of stars, ranging from 1 to 5.

We evaluate taste similarity between users that are assigned close-by profiles in the spectral embedding. We do this as follows: We select a set of 2000 users and a set of 2000 movies from the Netflix data set, such that the selected users have given roughly the same number of ratings within the selected movie set. The presence of a rating is viewed in this setting as a sign that the user has viewed that particular movie, and is therefore considered as a binary form of appreciation (the lack of a rating denoting a potential lack of interest for that content). Subsequently, for each user we hide the rating of one content at random. Using the remaining observed ratings, we build a sparse observed 0 – 1 rating matrix S , and we compute the spectral profiles $\sqrt{N}z_u$ of the users. For each user, an ordered list of neighbours is implicitly defined, from the closest to the farthest one in the profile space. We compute over the set of users the average frequency of the occurrence of the following event: “a user at distance k has rated the hidden content”. The average is taken over the set of users. This “frequency of agreement” reflects the taste proximity of users.

In Figure 1 we plot this “frequency of agreement” for different values of the dimensionality L of the embedding. Content popularity (i.e. the fraction of users having rated it) ranges from 0.45% to 5.3%, and the average popularity of content is 2.13%. We show this value on our plots for comparison. In Figure 1(a) we vary the dimension from 2 to 150 and plot the average frequency of agreement with the nearest neighbour. We notice that there is a peak in this average frequency around roughly 30 dimensions. Subsequently the plotted frequency decays slowly. In Figure 1(b) we plot the average frequency of agreement for 2, 30 and 150 dimensions for the 100 closest neighbours in the profile space. We conclude that an embedding of rank 30 is appropriate to characterise user taste for the selected users in the trace.



(a) Average frequency of agreement with nearest neighbour



(b) Average frequency of agreement with neighbours

Figure 1: Profiling Proximity

It is important to observe that the frequency decays with distance in the profile space. This indicates that spectral profiling manages to capture user taste proximity.

4 Oja's Algorithm and Beyond

We now propose a method for extracting the eigenvectors of the adjacency matrix of a graph in a distributed manner. Eigenvector extraction is the object of Oja's algorithm [14], the basic version of which is refined by Borkar and Meyn in [2].

Consider a sequence of symmetric square random matrices $(A_k \in \mathbb{R}^{N \times N})_k$ of common finite mean $A \in \mathbb{R}^{N \times N}$. Oja and Karhunen [14] proposed the following stochastic approximation algorithm for determining the s top eigenvectors of A :

$$\tilde{X}_k = X_{k-1} + A_k X_{k-1} \Gamma_k, \quad (14)$$

$$X_k = \tilde{X}_k R_k^{-1}, \quad (15)$$

where X_k is an $N \times L$ matrix, Γ_k is a diagonal matrix of gains and R_k^{-1} is a matrix achieving the orthonormalisation of the columns of \tilde{X}_k . They prove almost sure convergence under typical assumptions on the sequence of gains, assuming unit multiplicity of the top s eigenvalues, probability density uniformly bounded away from 0 for each of the A_k , and almost sure boundedness and symmetry of the A_k , as well as statistical independence.

A simpler single-step form of the algorithm was also proposed by Oja and Karhunen. The convergence of a slightly modified version thereof was showed by Borkar and Meyn in [2], who introduced the additional factor $\frac{1}{1 + \text{Tr}(X_k X_k')}$ and added the additional i.i.d. $\mathcal{N}(0, I)$ noise sequence ξ_k :

$$X_k - X_{k-1} = \frac{a_k}{1 + \text{Tr}(X_k X_k')} [(I - X_{k-1} X_{k-1}') A_{k-1} X_{k-1} + \xi_k].$$

Here (a_k) is an almost typical gain sequence:

$$\sum_k a_k = \infty, \quad \sum_k a_k^2 < \infty, \quad \sup_k \frac{\sum_{n \geq k} a_n^2}{a_k} < \infty.$$

None of these approaches are distributed however. Our contribution consists precisely in augmenting these methods to make them distributed.

4.1 A Method for Distributed Spectral Profiling

In Section 3 we have demonstrated the benefits of spectral profiling. We have seen that given a set of N users, it is sufficient that each of them contact on average only $\Omega(\log(N))$ other users at random and determine similarity with them to essentially manage to characterise the profiles of everyone, by application of the spectral transformation.

In this section we consider a (sparse) graph obtained via such a probing process. We develop message-passing algorithms that enable all users to individually compute their spectral profile, while allowing communication only between pairs of neighbours. The algorithms we propose in this section only require the connectivity property of the graph (and hence no specific bound on neighbourhood sizes). However, the scenario of a sparse graph is most appealing, as it fits very well with the model we introduced in the previous section.

Let us thus consider a network described by an undirected graph $\mathcal{G} = (\mathcal{U}, \mathcal{E})$, where \mathcal{U} is the set of N nodes (also designated interchangeably by “users”), and \mathcal{E} is the set of edges connecting these nodes. We denote by \mathcal{N}_u the set of neighbours of u , and also write $u \sim v$ to indicate that two nodes u, v are neighbours. Nodes that share a link will compute their “similarity” value. The user similarity values computed for peers connected by an edge define an adjacency matrix, which we call A . In this section we show that nodes are able to compute individually via message passing their coordinates in an L -dimensional profile space. These coordinates form a collection of L linearly independent vectors which span the vector space generated by the L eigenvectors corresponding to the top L largest magnitude eigenvalues of A .

In the context of Section 3.1, we can think of the graph \mathcal{G} as a random $G(N, p)$ graph, where p is the probing probability. Since in the model we have introduced the eigenvectors are quasi-constant on each of the user profiles, a linear combination thereof will have the same property. Hence, the results of Theorem 1 hold for such a collection of vectors, for a modified version of condition (2c). More specifically, consider a full rank $L \times L$ matrix $W = (w_{k\ell})$ of linear coefficients, and the matrix $X = (x_1, \dots, x_L)$ of eigenvectors of the adjacency matrix A . Results in Section 3 hold for spectral profile vectors given by XW , under the following condition:

The matrix of normalised eigenvectors $Y = (y_1, \dots, y_L)$ under the α -norm of matrix $M = B\text{diag}(\alpha)$ satisfy

$$t'_k \neq t'_\ell, 1 \leq k \neq \ell \leq K,$$

$$\text{where } t'_k := (Yw_1(k), \dots, Yw_L(k)), \quad (2c').$$

For presentation ease, in the following we assume that A is positive semidefinite.

We now describe our proposed method. Given some fixed number L of target eigenvectors to be extracted, each user u maintains at all time t three sets of variables:

1. An L -dimensional row vector $X_u(t)$, the sought-for eigenvector coordinates;
2. An $L \times L$ matrix $\Phi_u(t)$ and a scalar $\Psi_u(t)$, both playing an auxiliary role in the calculation.

For ease of presentation, we assume slotted time $t = 0, 1, 2, \dots$, and synchronous updates at all peers. Asynchronous versions will be described and tested in the next section.

Our algorithm then takes the following form:

$$X_u(t+1) - X_u(t) = \frac{a(t)}{Y_u(t)} \left[\sum_{v \sim u} A_{vu} X_v(t) - N X_u(t) \Phi_u(t) + \xi_u(t+1) \right]. \quad (16)$$

In the above, $a(t)$ is a gain parameter to be specified, $\xi_u(t+1)$ is a noise term deliberately introduced by

user u , and the denominator $Y_u(t)$ is taken equal to

$$Y_u(t) = \max \left(1, |\Psi_u(t)|, \frac{1}{NL^2} \sum_{k,\ell=1}^L |(\Phi_u(t))_{k\ell}| \right). \quad (17)$$

It is readily seen that the update (16) can be computed locally at u , solely relying on variables local to node u and inputs $X_v(t)$ from u 's neighbours $v \in \mathcal{N}_u$. The same is also true for the updates of variables Φ_u and Ψ_u , which take the forms:

$$\Phi_u(t+1) = \Phi_u(t) + b(t) \sum_{v \sim u} (\Phi_v(t) - \Phi_u(t)) + f_u(t+1) - f_u(t), \quad (18)$$

where $b(t)$ is a gain parameter, $f_u(t)$ is a $L \times L$ matrix, specified by

$$f_u(t) = X'_u(t) \sum_{v \sim u} A_{uv} X_v(t), \quad (19)$$

and

$$\Psi_u(t+1) = \Psi_u(t) + b(t) \sum_{v \sim u} (\Psi_v(t) - \Psi_u(t)) + g_u(t+1) - g_u(t), \quad (20)$$

where $g_u(t)$ is a scalar, specified by

$$g_u(t) = X_u(t) X'_u(t). \quad (21)$$

Before stating the main result of this section, we introduce the technical conditions that will be required from the gain sequences $a(t)$, $b(t)$:

$$a(t), b(t) \in [0, 1], \quad t > 0, \quad (22a)$$

$$\sum_{t>0} a(t) = \sum_{t>0} b(t) = +\infty, \quad (22b)$$

$$\sum_{t>0} a(t)^2 < +\infty, \quad \lim_{t \rightarrow \infty} b(t) = 0, \quad (22c)$$

$$\lim_{t \rightarrow \infty} \frac{a(t)}{b(t)} e^{K \sum_{s=1}^t a(s)} = 0, \quad K > 0 \quad (22d)$$

Note that these conditions are satisfied for instance upon taking $a(t) = 1/(t \log(t))$, and $b(t) = t^{-2/3}$. Indeed, with this choice for $a(t)$, it is readily seen that

$$\sum_{s=1}^t a(s) \sim \log(\log(t)),$$

and (22b) follows. In addition, the quantity in (22d) then reads

$$\frac{1}{t^{1/3} \log(t)} e^{K(1+o(1)) \log(\log(t))} \leq \frac{(\log(t))^{2K-1}}{t^{1/3}},$$

where we have used the upper bound of 1 on the term $o(1)$, and property (22d) readily follows.

We are now in a position to state this section's main result:

Theorem 3. *Assume that the gains $a(t)$, $b(t)$ verify the conditions (22). Assume further that the noise terms $\xi_u(t)$ are i.i.d, white Gaussian noise. Assume finally that the overlay graph over which peers communicate is connected. Then the distributed updating algorithm (16–21) verifies the following property. With probability 1, the columns of $X(t) := (X_u(t))_{u \in \mathcal{U}}$ converge to a collection of L orthonormal vectors spanning the vector space associated with the L largest eigenvalues of the weighted adjacency matrix A .*

The proof of the theorem is given in Section F. In what follows, we provide some background and intuition for it.

Consider first the main equation, (16). If we ignore the denominator $Y_u(t)$, the noise term $\xi_u(t+1)$, and replace the term $NX_u(t)\Phi_u(t)$ by $\sum_v X_u(t)f_v(t)$, where $f(t)$ is as given in (19), this equation reads, written in matrix form:

$$X(t+1) - X(t) = a(t) [AX(t) - X(t)X'(t)AX(t)].$$

This is in fact the celebrated Oja algorithm, proposed by Oja and Karhunen [14] to extract precisely the eigenvectors of the largest eigenvalues of A . Oja's algorithm is subject to some stability issues, that Borkar and Meyn [2] proposed to fix by scaling down the right-hand side of the previous equation by some factor $Z(t) = 1 + \sum_{u,k} X_{u,k}^2(t)$, and by adding an extra noise term $\xi(t+1)$ in the bracket in the right-hand side. Thus, the update rule they considered reads:

$$X(t+1) - X(t) = \frac{a(t)}{Z(t)} [AX(t) - X(t)X'(t)AX(t) + \xi(t+1)], \quad (23)$$

and is proved in ([2]) to converge with probability 1 to the desired eigenvectors, under assumptions (22b), (22c) on the gains $a(t)$, and similar conditions on the noise $\xi(t)$ as in our theorem.

However, algorithm (23) does not lend itself to a distributed implementation, since neither of the two terms $X'(t)AX(t)$ or $Z(t)$ can be computed locally by the users.

To solve this issue, we introduce the auxiliary local variables Φ_u, Ψ_u . The dynamics (18–20) according to which they evolve is best understood by setting to zero the input terms $f_u(t+1) - f_u(t)$ and $g_u(t+1) - g_u(t)$ in the right-hand side. It then becomes apparent that these dynamics perform local averaging (also known as gossiping in [3]). Thus these eventually converge to a state where all variables $\Phi_u(t)$ coincide with the average $(1/N) \sum_v \Phi_v(0)$ of the original entries.

We can now provide a heuristic argument for the theorem. On a fast time scale, characterised by the gain parameters $b(t)$, the gossiping dynamics converge to almost constant vectors, with

$$\begin{aligned} \Phi_u(t) &\equiv \frac{1}{N} \sum_v f_v(t), & u \in \mathcal{U}, \\ \Psi_u(t) &\equiv \frac{1}{N} \sum_v g_v(t) & u \in \mathcal{U}. \end{aligned}$$

Then on a slower time scale dictated by the gain parameters $a(t)$, the variables of interest $X_u(t)$ follow dynamics very close to (23). Indeed, the auxiliary parameters Φ_u, Ψ_u track accurately the desired terms $X'(t)AX(t)$ and $Z(t)$ respectively.

A couple of remarks are in order. The stabilisation by the scaling factor $Z(t)$ in (23) seems insufficient in the presence of the additional dynamics (18,20). This leads us to introduce our alternative stabilisation via $Y_u(t)$ in (17). Also, in problems with dynamics at two time scales a common assumption on the gain parameters is that $a(t)/b(t) \rightarrow 0$. In the present case, a stronger form of time scale separation (namely, condition (22d)) is needed, to prevent reinforcing instabilities between the two dynamics.

4.2 Evaluation of an Asynchronous Version

In this section we present numerical evaluations on synthetic data. We exhibit convergence of an asynchronous version of the distributed coordinate assignment scheme presented in the previous section on synthetic data generated according to the model presented in Section 3.1.

In Section 4.1 we showed that the distributed algorithm (16–18) converges almost surely towards L linearly independent vectors spanning the vector space generated by the eigenvectors corresponding to the top L magnitude eigenvalues of the adjacency matrix A .

In the following, we evaluate the asynchronous version of the algorithm. In this setting, each node u keeps track of its own coordinates X_u as well as the gossiped variables Φ_u and Ψ_u . However, instead

Algorithm 1 Distributed profiling algorithm

NODE(u)::UPDATE-LOCAL() at rate λ Local Variables: $X_u, \Pi_u, w_u, X_u^0, \Pi_u^0, \Psi_u, \Phi_u$

```
1: if  $w_u = 1$  then
2:   Retrieve partial product vectors  $\Pi_v$  from  $v \in \mathcal{N}_u$ 
3:    $X_u := X_u + \gamma \frac{\sum_{v \sim u} A_{uv} \Pi_v - N X_u \Phi_u}{Y_u(\Psi_u, \Phi_u)}$ 
4: else
5:   Retrieve vectors  $X_v$  from  $v \in \mathcal{N}_u$ 
6:    $\Pi_u := \sum_{v \sim u} A_{uv} X_v$ 
7: end if
8:  $w_u := 1 - w_u$ 
```

LINK(u, v)::GOSSIP() at rate μ

```
1: Retrieve local variables at  $u$  and  $v$ 
2: for  $(H, h)$  in  $\{(\Phi, f^{(2)}), (\Psi, g)\}$  do
3:    $\alpha := \frac{H_u + H_v}{2}$ 
4:   for  $i$  in  $\{u, v\}$  do
5:      $\delta_i := h_i(X_i, \Pi_i) - h_i(X_i^0, \Pi_i^0)$ 
6:      $H_i := \alpha + \delta_i$ 
7:   end for
8: end for
9:  $X_u^0 := X_u, \Pi_u^0 := \Pi_u, X_v^0 := X_v, \Pi_v^0 := \Pi_v$ 
```

of explicitly imposing a timescale separation via gains $a(t)$ and $b(t)$ while enforcing a synchronised evolution of all the quantities, we impose distinct rates at which the updates are performed. Namely, the coordinate updates (16) are performed independently according to Poisson processes of rate λ , while gossiping (17-18) is performed independently “pairwise” according to Poisson processes of rate $\mu \gg \lambda$. By pairwise we mean that a pair of nodes $(u, v) \in \mathcal{E}$ will exchange and update their values for Φ and Ψ at rate μ similarly to the randomised gossiping technique from [3].

Furthermore, we replace the adjacency matrix by its square A^2 . We choose to do so, since the latter is positive semidefinite, has the same eigenvectors as A , and a spectrum composed of the squared eigenvalues of A . Implicitly, the eigenvectors are ordered according to the magnitude of the eigenvalues of A , instead of their actual values. In turn, this modification alters the function f_u from (19), which becomes:

$$f_u^{(2)} = \left(\sum_{v \sim u} A_{uv} X_v \right)' \left(\sum_{v \sim u} A_{uv} X_v \right). \quad (24)$$

The algorithm executed at each node is summarised in Algorithm 1.

Since sparsity is not preserved by taking the square of A , we cannot simply use A^2 as a new adjacency matrix. Thus, we need to compute products $A^2 X$ specifically: For some node u , every other call to the UPDATE-LOCAL() procedure computes the partial products $\Pi_u = \sum_{v \sim u} A_{uv} X_v$ (State 6 in the Algorithm). Subsequently, the neighbours’ partial product vectors Π_v are used in the coordinate update procedure (16) at State 3 and for the gossiping of $f_u^{(2)} = \Pi_u' \Pi_u$. Each node additionally stores its previous vectors X_u^0 and Π_u^0 for use in the GOSSIP() procedure.

The only piece of global information required is the number of users in the system (or an approximation thereof). In the UPDATE-LOCAL() procedure, we used a fixed gain γ . The noise component ξ is omitted in the algorithm. Even so, noise is intrinsic to the algorithm as it is introduced by both the gossip averaging and by the fact that exchanges are asynchronous.




Figure 2: Pure and noisy user profiles

For evaluation of convergence, we used a synthetic data trace generated according to our model from Section 3.1. This trace considers 2200 users clustered in 4 classes of sizes 200, 500, 600 and 900. The probability matrix pB characterising the classes is

$$pB = \frac{1}{100} \begin{pmatrix} 0.5 & 1.5 & 2.0 & 1.0 \\ 1.5 & 0.55 & 1.0 & 2.0 \\ 2.0 & 1.0 & 0.45 & 4.0 \\ 1.0 & 2.0 & 4.0 & 0.5 \end{pmatrix}.$$

We consider an adjacency matrix A generated using the aforementioned parameters. This matrix is sparse (the average node degree is 40). For visualisation ease, we consider the case $L = 2$ (i.e. the profile space is a 2-dimensional plane). In Figure 2 we plot the first two eigenvectors of the expected adjacency matrix \bar{A} . They are constant on each of the 4 classes, hence the plot is constituted of 4 points, which are depicted as the four black squares. They represent the pure user profiles which characterise each of the four classes. Additionally, we plot the first two eigenvectors of the “noisy” adjacency matrix A , with elements belonging to a specific class marked with the same symbol. We have the visual confirmation that despite the sparsity of matrix A , the noisy profiles are grouped around the pure profiles.

We choose a gain $\gamma = 0.001$, an update rate $\lambda = 0.2$ and a gossip rate $\mu = 10$. We initialise the algorithm at a random state. In Figure 4.2 we plot the time evolution of the proportion of the mass of the two coordinate vectors $X_{\cdot 1}$ and $X_{\cdot 2}$ (aggregated across users) that falls on the space orthogonal to the 2-dimensional eigenspace generated by the first two eigenvectors of matrix A^2 . Additionally, we plot the scalar product of the two coordinate vectors. After roughly 400 time units, we observe convergence towards orthogonal vectors spanning the desired eigenspace.

5 Conclusions

In this paper we addressed the problem of distributed user profiling and recommendation.

We first showed that spectral techniques constitute an appealing approach, and obtained novel results on their efficiency, thereby improving upon previous literature on the subject of spectral clustering. We showed that for a low-rank probabilistic model of user taste, a simple distributed algorithm based on local votes in the profile space asymptotically achieves accurate prediction of user preference.




Figure 3: Convergence of the asynchronous distributed algorithm

We developed techniques for computing eigenvectors in a distributed manner. Our solution combines ideas from Oja’s algorithm with gossiping algorithms. From a theoretical standpoint, it essentially relies on a special form of multiple time scale stochastic approximation. The resulting technique may have other applications besides user profiling.

Finally we evaluated our proposed methods on synthetic and actual data traces. We thereby validated our analysis in observing convergence to the desired eigenvectors. We could further show, based on the Netflix prize dataset, that accurate recommendations could be made at limited communication cost based on our spectral embedding.

Several research directions can be envisioned to take this work further. One intriguing problem concerns privacy. While our methods do not rely on direct exchange of sensitive private information, they may nevertheless lead to private information leakage. A distributed solution avoiding the issue is yet unavailable.

Other directions concern the fine tuning of the methods. The issue of analytical selection of the number of eigenvectors has not been addressed here. The recent work of Shi et al. [15] could be an appealing solution.

References

- [1] Xavier Amatriain, Josep Pujol, and Nuria Oliver. I like it... i like it not: Evaluating user ratings noise in recommender systems. volume 5535 of *Lecture Notes in Computer Science*, pages 247–258. Springer Berlin / Heidelberg, 2009. 10.1007/978-3-642-02247-0_24.
- [2] V. Borkar and S. P. Meyn. Oja’s algorithm for graph clustering and markov spectral decomposition. In *ValueTools ’08: Proceedings of the 3rd International Conference on Performance Evaluation Methodologies and Tools*, pages 1–7, ICST, Brussels, Belgium, Belgium, 2008. ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering).
- [3] Stephen Boyd, Arpita Ghosh, Balaji Prabhakar, and Devavrat Shah. Randomized gossip algorithms. *IEEE/ACM Trans. Netw.*, 14(SI):2508–2530, 2006.
- [4] Anirban Dasgupta, John Hopcroft, Ravi Kannan, and Pradipta Mitra. Spectral clustering by recursive partitioning. In *ESA’06: Proceedings of the 14th conference on Annual European Symposium*, pages 256–267, London, UK, 2006. Springer-Verlag.
- [5] Uriel Feige and Eran Ofek. Spectral techniques applied to sparse random graphs. *Random Struct. Algorithms*, 27(2):251–275, 2005.
- [6] David Kempe and Frank McSherry. A decentralized algorithm for spectral analysis. In *STOC*, pages 561–568, 2004.
- [7] Raghunandan H. Keshavan, Sewoong Oh, and Andrea Montanari. Matrix completion from a few entries. *CoRR*, abs/0901.3150, 2009.
- [8] Satish Babu Korada, Andrea Montanari, and Sewoong Oh. Gossip pca. In *Proceedings of the ACM SIGMETRICS joint international conference on Measurement and modeling of computer systems*, SIGMETRICS ’11, pages 209–220, New York, NY, USA, 2011. ACM.
- [9] Yehuda Koren. The bellkor solution to the netflix grand prize, 2009.
- [10] F. McSherry. Spectral partitioning of random graphs. *Proceedings FOCS*, pages 529–537, 2001.
- [11] A. Müller and D. Stoyan. *Comparison Methods for Stochastic Models and Risks*. J. Wiley and Sons, 2002.
- [12] Netflix. Netflix prize. <http://www.netflixprize.com>
- [13] Andrew Y. Ng, Michael I. Jordan, and Yair Weiss. On spectral clustering: Analysis and an algorithm. In *Advances in Neural Information Processing Systems 14*, pages 849–856. MIT Press, 2001.
- [14] E. Oja and J. Karhunen. On stochastic approximation of the eigenvectors and eigenvalues of the expectation of a random matrix. *Journal of Math. An. and App.*, 106(1), 1985.
- [15] Tao Shi, Mikhail Belkin, and Bin Yu. Data spectroscopy: learning mixture models using eigenspaces of convolution operators. In *ICML*, pages 936–943, 2008.
- [16] G. W. Stewart. On the early history of the singular value decomposition. *SIAM Review*, 35(4):551–566, 1993.
- [17] Dan-Cristian Tomozei and Laurent Massoulié. Distributed user profiling via spectral methods. *SIGMETRICS Perform. Eval. Rev.*, 38:383–384, June 2010.

A Proof of Lemma 1

Consider an eigenvalue $\bar{\lambda}$, $|\bar{\lambda}| > 0$ of \bar{A} and a corresponding eigenvector \bar{x} . For $1 \leq u \leq N$,

$$(\bar{A}\bar{x})(u) = \frac{\omega}{N} \sum_{v=1}^N b_{k(u)k(v)} \bar{x}(v) = \frac{\omega}{N} \sum_{\ell=1}^K b_{k(u)\ell} \sum_{v:k(v)=\ell} \bar{x}(v) = \bar{\lambda} \bar{x}(u).$$

For a large enough N , we have that $|C_k| > 1$, for all k . This is true, since the size of each class grows linearly with N . Then for all k and for all $u, u' \in C_k$, with $u' \neq u$ it follows that $\bar{x}(u') = \bar{x}(u)$.

Denote the value of $\bar{x}(u)$ by $y(k(u))$. Then, for all $1 \leq \ell \leq K$,

$$\sum_{\ell'=1}^K \alpha_{\ell'} b_{\ell\ell'} y(\ell') = \frac{\bar{\lambda}}{\omega} y(\ell).$$

Thus y is an eigenvector of the $K \times K$ matrix $M = B \text{diag}(\alpha)$ corresponding to its eigenvalue $\frac{\bar{\lambda}}{\omega}$. Since M is a constant matrix, its eigenvalues are also constants. Hence it must be that there exists a constant c such that $\bar{\lambda} = c\omega = \Theta(\omega)$. By Condition (2b) we conclude that the top L magnitude eigenvalues of \bar{A} have distinct absolute values.

Finally, we have that

$$\|\bar{x}\|_2 = \sqrt{N} \|y\|_\alpha.$$

Thus, if we consider a vector y normalised under the α -norm, the components of the normalised eigenvector $\frac{\bar{x}}{\|\bar{x}\|}$ are given by (3). \square

Note that a consequence of this result is that \bar{A} can have a number of non-zero eigenvalues that is at least L (by Condition (2a)) and at most K (i.e. the maximum number of eigenvalues of matrix M).

B Proof of Lemma 2

We start by establishing a simple result of stochastic dominance.

By definition a random variable X is dominated for the convex ordering by another random variable Y (written as $X \leq_{cx} Y$) if for any convex function $f : [c, d] \rightarrow \mathbb{R}$ such that $\mathbb{E}f(X)$ and $\mathbb{E}f(Y)$ exist, we have $\mathbb{E}f(X) \leq \mathbb{E}f(Y)$.

It is known that (see for instance [11], Theorem 1.5.20): If $X \leq_{cx} Y$, then there exist \hat{X} and \hat{Y} with the same distributions as the original variables, but which are such that $\mathbb{E}(\hat{Y}|\hat{X}) = \hat{X}$.

Another result found in [11] states that for some closed interval $[c, d]$, if $X : \Omega \rightarrow [c, d]$ and $Y : \Omega \rightarrow [c, d]$ are two random variables such that $\mathbb{E}X = \mathbb{E}Y$, then $X \leq_{cx} Y$.

In this latter setting, we wish to establish a variant of the former result. Namely, for some closed interval $[c, d]$, if $X : \Omega \rightarrow [c, d]$, we wish to construct a random variable $\hat{Y} : \Omega \rightarrow \{c, d\}$ supported on the extremities $\{c, d\}$ of the interval, such that $\mathbb{E}X = \mathbb{E}\hat{Y}$ and $\mathbb{E}(\hat{Y}|X) = X$. To achieve this, pick a uniformly distributed random variable $U \sim \mathcal{U}[0, 1]$ independent of X . We define the random variable $\hat{Y} : \Omega \rightarrow \{c, d\}$ as $\hat{Y} = F(X, U)$, where

$$F(x, u) = d - \mathbb{1}_{\{x \leq u(d-c)+c\}}(d-c).$$

Then

$$\mathbb{E}(\hat{Y}|X) = \mathbb{E}(F(X, U)|X) = d - \frac{d-X}{d-c}(d-c) = X.$$

Since \hat{Y} can possibly take only two values, we can compute the corresponding probabilities:

$$\mathbb{P}(\hat{Y} = d) = \int_c^d P_X(dx) \int_0^{\frac{x-c}{d-c}} du = \frac{\mathbb{E}X - c}{d - c} \text{ and } \mathbb{P}(\hat{Y} = c) = \frac{d - \mathbb{E}X}{d - c}, \quad (25)$$

and hence $\mathbb{E}X = \mathbb{E}\hat{Y}$.

Let us now proceed with the proof of the Lemma.

Denote by $Q := A - \mathbb{E}A$. Denote further by $p_{max} = \max_{ij} p_{ij}$ and by $p_{min} = \min_{ij} p_{ij}$. Then the elements of Q all belong to the interval $[-p_{max}, 1 - p_{min}]$. For a large enough N such that $p_{max} < 1$, there exist $\alpha = \frac{1-p_{min}}{1-p_{max}} > 1$ and $p = p_{max}$, such that $[-p_{max}, 1 - p_{min}] \subset \alpha[-p, 1 - p]$.

Consider a symmetric matrix U of independent uniformly distributed random variables $\{U_{ij} = U_{ji} \sim \mathcal{U}[0, 1]\}_{i < j}$. Then for each entry Q_{ij} of Q , such that $i < j$, which we regard as a random variable with values in the interval $\alpha[-p, 1 - p]$, we construct the random variable $Z_{ij} = F(Q_{ij}, U_{ij})$ which has the desired property written in matrix form

$$\mathbb{E}(Z|Q) = Q. \quad (26)$$

Note that the entries of Z are mutually independent by construction. Furthermore, the random variables defined as $\{Y_{ij} = \alpha^{-1}Z_{ij} + p, Y_{ji} = Y_{ij}\}_{i < j}$ are mutually independent Bernoulli random variables of parameter p and hence form the adjacency matrix Y of an Erdos-Renyi graph of parameters (N, p) . Denote $\bar{Y} := \mathbb{E}Y = p(ee' - I)$, where e is the all-ones column vector and e' denotes transposition.

Let us now prove that the spectral radius of $Z = \alpha(Y - \bar{Y})$ is upper bounded by $O(\sqrt{\omega})$ with high probability.

Since $\omega = \Omega(\log N)$, and the $\{Y_{ij}\}_{i < j}$ are mutually independent, we can apply the results from [5]. Let y be any vector of norm 1 and denote $u := \frac{1}{\sqrt{N}}e$. We can decompose y as follows: $y = ax + bu$, where $a^2 + b^2 = 1$, and x is a vector of norm 1 orthogonal to u , $x \perp u$.

$$|y'(Y - \bar{Y})y| \leq 2 \underbrace{|abx'(Y - \bar{Y})u|}_{T_1} + a^2 \underbrace{|x'(Y - \bar{Y})x|}_{T_2} + b^2 \underbrace{|u'(Y - \bar{Y})u|}_{T_3}.$$

Denote by δ_i the degree of node i and by $\bar{\delta} := \frac{\sum_i \delta_i}{N}$ the average degree. According to Lemma 2.2 from [5] and taking into account the fact that e is an eigenvector of \bar{Y} we have

$$T_1 = |x'Yu| \leq 2\sqrt{\bar{\delta}}, \text{ with probability } 1 - e^{-\Omega((N\omega)^{1/3})}.$$

By Theorem 2.5 and Claim 2.4 from [5], we have that for every constant $c_1 > 0$, there exists another constant $c_2 > 0$ such that:

$$|x'Yx| \leq c_2\sqrt{\omega}, \text{ with probability } 1 - N^{-c_1}. \quad (27)$$

We will restrict ourselves to constants $c_1 > 1$ for reasons that will become apparent later in the proof. Thus, we can bound the second term with probability $1 - N^{-c_1}$:

$$\begin{aligned} T_2 &\leq |x'Yx| + |x'\bar{Y}x| \stackrel{(27)}{\leq} O(\sqrt{\omega}) + \left| \sum_i x_i \sum_{j: i \neq j} px_j \right| \\ &= O(\sqrt{\omega}) + \left| -p \sum_i x_i^2 \right| = O(\sqrt{\omega}) + \Theta\left(\frac{\omega}{N}\right). \end{aligned}$$

Finally, using a Chernoff bound we find

$$T_3 = |\bar{\delta} - \omega| = O(\sqrt{\omega}), \text{ with prob. } 1 - e^{-\Omega(N)}.$$

Thus, $\rho(Z) = O(\sqrt{\omega})$, with probability $1 - N^{-c_1}$. Furthermore, there exists a constant $a > 0$ such that $\rho(Z) < aN$.

Let us now finally characterise the spectral radius of Q .

Using the fact that the spectral radius is a convex function and by Jensen's inequality, we get

$$\rho(Q) = \rho(\mathbb{E}[Z|Q]) \leq \mathbb{E}[\rho(Z)|Q]$$

We have a random variable $R := \rho(Z)$ supported on $[0, aN]$, such that $\mathbb{P}(R > t) \leq O(N^{-k})$ for $t = O(\sqrt{\omega})$ and we wish to deduce that the conditional expectation $S := \mathbb{E}(R|Q)$ is also upper bounded by $O(\sqrt{\omega})$ with high probability.

Since R and S have countable state spaces, it makes sense to consider

$$\beta(s) := \mathbb{P}(R > t | S = s).$$

Since $R : \Omega \rightarrow [0, aN]$, and since $\mathbb{E}(R|S) = S$, we have that $\beta(s)(aN) + (1 - \beta(s))t \geq s$, and thus, $\beta(s) \geq (s - t)/(aN - t)$. Denoting $\gamma := \mathbb{P}(S > t + 1)$, we have

$$\mathbb{P}(R > t) = E(\beta(S)) \geq E(\beta(S)\mathbb{1}_{\{S > t+1\}}) \geq \frac{t + 1 - t}{aN - t}\gamma = \frac{\gamma}{aN - t}.$$

Hence

$$\gamma = \mathbb{P}(S > t + 1) \leq (aN - t)\mathbb{P}(R > t) = (aN - t)O(N^{-c_1}) = o(1),$$

since we considered $c_1 > 1$. □

C Proof of Lemma 3

We will show the two claims (5) and (6) by induction. Denote (5) by \mathcal{P}_k and (6) by \mathcal{Q}_k .

We will begin by proving \mathcal{P}_1 and \mathcal{Q}_1 . Since we make extensive use of Lemma 2, it is implied that all inequalities in this proof hold with high probability in the sense of Lemma 2 (that is with probability $1 - e^{-\Omega(N)}$). All vectors are column vectors and we denote by x' the transpose of vector x . Furthermore, for two vectors x and y by $x \perp y$ we mean that their scalar product $x'y$ equals zero.

Using the variational characterisation of eigenvalues and Lemma 2, we get:

$$||\lambda_1| - |\bar{\lambda}_1|| \stackrel{\text{wlog}}{=} |\lambda_1| - |\bar{\lambda}_1| \leq |\lambda_1| - |x'_1 \bar{A} x_1| \leq |x'_1 (A - \bar{A}) x_1| \leq O(\sqrt{\omega}),$$

which proves \mathcal{P}_1 .

We denote the first eigenvector of A by $x_1 = a_1 \bar{x}_1 + b_1 \bar{y}_1$, where $a_1^2 + b_1^2 = 1$, \bar{x}_1 is the first eigenvector of \bar{A} and $\bar{x}_1 \perp \bar{y}_1$. Then, by making use of \mathcal{P}_1 , there exist positive constants θ_1 and θ_2 such that

$$|\bar{\lambda}_1| - \theta_1 \sqrt{\omega} \leq |\lambda_1| \leq a_1^2 |\bar{x}'_1 \bar{A} \bar{x}_1| + b_1^2 |\bar{y}'_1 \bar{A} \bar{y}_1| + \theta_2 \sqrt{\omega}.$$

We took into account the symmetry of \bar{A} and the fact that $\bar{A} \bar{x}_1 = \bar{\lambda}_1 \bar{x}_1$. By the Courant-Fischer theorem, we get the following inequality:

$$|\bar{\lambda}_1| - \theta \sqrt{\omega} \leq |\bar{\lambda}_1| - b_1^2 (|\bar{\lambda}_1| - |\bar{\lambda}_2|), \quad \theta > 0 \tag{28}$$

and since the top L eigenvalues of \bar{A} are distinct (by Condition (2b)), it holds that $0 < |\bar{\lambda}_1| - |\bar{\lambda}_2| = \Theta(\omega)$. We get that $b_1^2 \leq O(\frac{1}{\sqrt{\omega}})$, thus proving \mathcal{Q}_1 .

In order to generalise this result, we make use of the following simple lemma:

Lemma 6. Let x_1 and \bar{x}_1 be two non-orthogonal vectors of norm 1 such that

$$1 - (x'_1 \bar{x}_1)^2 \leq O\left(\frac{1}{\sqrt{\omega}}\right). \quad (29)$$

Then for all vectors x_2 of norm 1 such that $x_2 \perp x_1$, $(x'_2 \bar{x}_1)^2 \leq O\left(\frac{1}{\sqrt{\omega}}\right)$.

Proof. We have $x_1 = a\bar{x}_1 + b\bar{y}_1$, where $\bar{x}_1 \perp \bar{y}_1$, $\|\bar{y}_1\| = 1$ and $a^2 + b^2 = 1$. By hypothesis (29), $b^2 \leq O\left(\frac{1}{\sqrt{\omega}}\right)$. Thus,

$$x'_2 x_1 = ax'_2 \bar{x}_1 + bx'_2 \bar{y}_1 = 0,$$

and thus, since $a \neq 0$,

$$|x'_2 \bar{x}_1|^2 = \frac{b^2}{1 - b^2} |x'_2 \bar{y}_1|^2 \leq \theta \frac{1}{\sqrt{\omega}} \frac{\sqrt{\omega}}{\sqrt{\omega} - \theta},$$

where $\theta > 0$ is a constant. \square

We proceed by complete induction. We showed \mathcal{P}_1 and \mathcal{Q}_1 . Now assume \mathcal{P}_ℓ and \mathcal{Q}_ℓ are true for all $1 \leq \ell < k$. We wish to show \mathcal{P}_k and \mathcal{Q}_k .

Let us write $x_k = \alpha \bar{x}_k + \beta \bar{y} + \gamma \bar{z}$, where $\bar{y} \in \text{Span}\{\bar{x}_{k+1}, \bar{x}_{k+2}, \dots\}$ and $\bar{z} \in \text{Span}\{\bar{x}_1, \dots, \bar{x}_{k-1}\}$ and $\alpha^2 + \beta^2 + \gamma^2 = 1$. Lemma 6 and the induction hypothesis show that $\gamma^2 \leq O\left(\frac{1}{\sqrt{\omega}}\right)$. Since

$$|x'_k \bar{A} x_k| \leq (1 - \gamma^2) |\bar{\lambda}_k| + \gamma^2 |\bar{\lambda}_1| \leq |\bar{\lambda}_k| + O(\sqrt{\omega}),$$

we can conclude that

$$||\lambda_k| - |\bar{\lambda}_k|| \stackrel{\text{wlog}}{=} |\lambda_k| - |\bar{\lambda}_k| \leq |\lambda_k| - |x'_k \bar{A} x_k| + O(\sqrt{\omega}) \leq O(\sqrt{\omega}),$$

thus proving \mathcal{P}_k .

We have that,

$$\begin{aligned} |\lambda_k| &\leq |x'_k \bar{A} x_k| + |x'_k (A - \bar{A}) x_k| \leq \alpha^2 |\bar{\lambda}_k| + \beta^2 |\bar{\lambda}_{k+1}| + \gamma^2 |\bar{\lambda}_1| + O(\sqrt{\omega}) \\ &\leq |\bar{\lambda}_k| + \frac{a}{\sqrt{\omega}} (|\bar{\lambda}_1| - |\bar{\lambda}_k|) - \beta^2 (|\bar{\lambda}_k| - |\bar{\lambda}_{k+1}|) + \theta \sqrt{\omega}, \end{aligned}$$

where a and θ are positive constants. Without loss of generality we now assume that $||\lambda_k| - |\bar{\lambda}_k|| = |\bar{\lambda}_k| - |\lambda_k|$ and using \mathcal{P}_k we get

$$\beta^2 \leq \frac{\varphi \sqrt{\omega} + \frac{a}{\sqrt{\omega}} (|\bar{\lambda}_1| - |\bar{\lambda}_k|)}{|\bar{\lambda}_k| - |\bar{\lambda}_{k+1}|} \leq O\left(\frac{1}{\sqrt{\omega}}\right), \quad \varphi > 0$$

thus proving \mathcal{Q}_k . \square

D Proof of Lemma 4

Consider a non-zero singular value of \bar{S} , $\sigma > 0$ and corresponding left and right singular vectors \bar{x} and \bar{y} . Then we can write:

$$(\bar{S}' \bar{x})_i = \sum_{u=1}^N \bar{S}_{ui} \bar{x}_u = \sigma \bar{y}_i \stackrel{\text{not}}{=} \sigma h_{k'(i)}, \quad (\bar{S} \bar{y})_i = \sum_{i=1}^F \bar{S}_{ui} \bar{y}_i = \sigma \bar{x}_u \stackrel{\text{not}}{=} \sigma g_{k(u)},$$

since $\bar{S}_{ui} = pr_{k(u)k'(i)}$ depends only on the class of u , $k(u)$, and the class of i , $k'(i)$. Since \bar{y} is a right-singular vector of \bar{S} , it is also an eigenvector of $\bar{S}'\bar{S}$ corresponding to eigenvalue σ^2 , and thus after some simplification we have that

$$\frac{\sigma^2}{\gamma\omega^2}g = R\text{diag}(\beta)R'\text{diag}(\alpha)g,$$

that is the K -dimensional vector g is an eigenvector of the constant matrix $R\text{diag}(\beta)R'\text{diag}(\alpha)$ corresponding to the eigenvalue $\frac{\sigma^2}{\gamma\omega^2}$. Since the eigenvalues of a constant matrix are constant, it follows that there exists a constant λ such that $\sigma = \lambda\omega = \Theta(\omega)$. By Condition (7b) it follows that the largest L singular values of \bar{S} are distinct. Moreover, matrix $R\text{diag}(\beta)R'\text{diag}(\alpha)$ can have at most K distinct non-zero eigenvalues and hence the same holds for the singular values of \bar{S} .

Finally, we have that

$$\|\bar{x}\|_2 = \sqrt{N}\|g\|_\alpha.$$

Thus, if we consider a vector g normalised under the α -norm, the components of the normalised eigenvector $\frac{\bar{x}}{\|\bar{x}\|}$ are given by (3). \square

E Proof of Lemma 5

For $\bar{A} = \tau\bar{S}$ we denote by $\bar{\zeta}_k^+$ the normalised eigenvector corresponding to the eigenvalue $\bar{\sigma}_k$ and by $\bar{\zeta}_k^-$ the normalised eigenvector corresponding to the eigenvalue $-\bar{\sigma}_k$. We introduce similar notation for the eigenvectors of $A = \tau S$, namely ζ_k^+ and ζ_k^- . Then it holds that

$$\bar{\zeta}_k^+ = \frac{1}{\sqrt{2}} \begin{bmatrix} \bar{x}_k \\ \bar{y}_k \end{bmatrix}, \bar{\zeta}_k^- = \frac{1}{\sqrt{2}} \begin{bmatrix} \bar{x}_k \\ -\bar{y}_k \end{bmatrix}, \zeta_k^+ = \frac{1}{\sqrt{2}} \begin{bmatrix} x_k \\ y_k \end{bmatrix}, \zeta_k^- = \frac{1}{\sqrt{2}} \begin{bmatrix} x_k \\ -y_k \end{bmatrix}.$$

By Condition (7b), we can apply a slightly modified Lemma 3 to matrix A . The only modification we need to make is to change the considered ordering of the eigenvalues – instead ordering them by largest magnitude, we order them decreasingly by value. Since we can apply Lemma 2 to $A - \bar{A}$, it is straightforward to see that the proof also holds in this setting.

We have thus that (10) holds and that with high probability

$$\sin(\widehat{\zeta_k^+, \bar{\zeta}_k^+}) \leq O(\omega^{-1/4}).$$

To see that (11) holds as well, we write

$$\begin{aligned} 1 - \langle \zeta_k^+, \bar{\zeta}_k^+ \rangle^2 &= 1 - \frac{1}{4}(\langle x_k, \bar{x}_k \rangle + \langle y_k, \bar{y}_k \rangle)^2 \leq O(\omega^{-1/2}), \\ -\langle \zeta_k^+, \bar{\zeta}_k^- \rangle^2 &= -\frac{1}{4}(\langle x_k, \bar{x}_k \rangle - \langle y_k, \bar{y}_k \rangle)^2 \leq 0. \end{aligned}$$

By summing the two expressions we get that:

$$\frac{1}{2}(1 - \langle x_k, \bar{x}_k \rangle^2) + \frac{1}{2}(1 - \langle y_k, \bar{y}_k \rangle^2) \leq O(\omega^{-1/2}),$$

hence the conclusion. \square

F Proof of Theorem 3

The main ingredient in the proof is the following

Lemma 7. *The update equation (16) is such that for all $t > 0$,*

$$\|X(t+1)\| \leq e^{K \sum_{s=1}^t a(s)} (\|X(1)\| + M) \quad (30)$$

for some positive constant K and some almost surely finite random variable M .

Proof. Rewrite Equation (16) in matrix form as

$$X(t+1) - X(t) = a(t) [F((X, \Phi, \Psi)(t)) + ND^{-1}(t)\xi(t+1)]$$

for some suitable function F , and where $D(t)$ denotes the $N \times N$ diagonal matrix with diagonal entries $Y_u(t)$. By the specific form of the terms $Y_u(t)$, and their role in the function F , it is readily seen that the latter verifies

$$\|F(X, \Phi, \Psi)\| \leq K_1 \|X\|$$

for some suitable constant K_1 . This readily implies that

$$\|X(t+1)\| \leq (1 + K_1 a(t)) \|X(t)\| + a(t) \eta(t),$$

where the $\eta(t) = N \|\xi(t+1)\|$ are iid. By induction, one then establishes that

$$\|X(t+1)\| \leq \prod_{s=1}^t (1 + K_1 a(s)) \left[\|X(1)\| + \sum_{s=1}^t a(s) \eta(s) \right].$$

Denote now $\bar{\eta}$ the expectation of $\eta(s)$, and let $M(t) := \sum_{s=1}^t a(s) [\eta(s) - \bar{\eta}]$. It is readily seen that $M(t)$ is a uniformly integrable martingale, and hence the supremum $\sup_{s \geq 0} M(s)$ is almost surely finite; denote it by \hat{M} . It then follows from the above equation that

$$\|X(t+1)\| \leq \prod_{s=1}^t (1 + K_1 a(s)) \left[\|X(1)\| + \sum_{s=1}^t \bar{\eta} a(s) + \hat{M} \right].$$

Using the elementary inequality $1 + x \leq e^x$, one deduces:

$$\|X(t+1)\| \leq e^{\sum_{s=1}^t K_1 a(s)} \left[\|X(1)\| + e^{\sum_{s=1}^t \bar{\eta} a(s)} + \hat{M} \right].$$

The result (30) then follows by setting $K = K_1 + \bar{\eta}$, and $M = 1 + \max(0, \hat{M})$. □

This Lemma is now used to establish the following result:

Lemma 8. *The auxiliary variables $\Phi_u(t)$, $\Psi_u(t)$ verify*

$$\begin{aligned} \lim_{t \rightarrow \infty} |N\Phi_u(t) - \sum_v f_v(t)| &= 0, \\ \lim_{t \rightarrow \infty} |N\Psi_u(t) - \sum_v g_v(t)| &= 0, \end{aligned} \quad (31)$$

i.e. asymptotically, these quantities do track accurately their intended targets.

Proof. We shall only consider the case of $\Psi_u(t)$, the other one being entirely similar. Rewrite the update rule (18) in matrix form as

$$\Psi(t+1) = (I - b(t)\Delta)\Psi(t) + g(t+1) - g(t), \quad (32)$$

where Δ is the so-called Laplacian matrix of the overlay graph: $\Delta_{uu} = |\mathcal{N}_u|$, $\Delta_{uv} = -\mathbb{1}_{u \sim v}$ for $u \neq v$. Recall that the Laplacian Δ is positive semi-definite, with eigenvector $e = (1, \dots, 1)'$ associated to the

eigenvalue 0. Also, when the overlay graph is connected, all other eigenvectors are associated with strictly positive eigenvalues $\lambda > 0$.

Note then that, because $g(t+1)$ and $g(t)$ have the same quadratic dependence on $X(t+1)$ and $X(t)$ respectively, it follows from Lemma 7 and the bound on function F therein that

$$\|g(t+1) - g(t)\| \leq a(t)K'e^{2K\sum_{s=1}^t a(s)}$$

for some suitable fixed finite term K' . By our Condition (22d) this is asymptotically negligible compared to $b(t)$. Let z be an eigenvector of Δ associated with a positive eigenvalue λ . Denote $z(t)$ the scalar product $z'\Psi(t)$. One then deduces from (32) the equation

$$z(t+1) = (1 - \lambda b(t))z(t) + b(t)\epsilon(t),$$

where $\lim_{t \rightarrow \infty} \epsilon(t) = 0$. By induction, one can deduce from the previous expression the identity

$$z(t+1) = z(1) \prod_{s=1}^t (1 - \lambda b(s)) + \sum_{s=1}^t \epsilon(s)b(s) \prod_{\sigma=s+1}^t (1 - \lambda b(\sigma)).$$

Elementary analysis can then be used to deduce from this last display, assumptions (22b),(22c) on gains $b(t)$, and convergence of $\epsilon(t)$ to 0 that $z(t)$ also converges to 0.

Indeed, take any fixed $\varepsilon > 0$. Since $\epsilon(t) \rightarrow 0$, there exists some t_0 such that $\epsilon(t) < \frac{\varepsilon\lambda}{3}$, for all $t \geq t_0$. Then, we can write:

$$z(t+1) \leq z(1) \prod_{s=1}^t (1 - \lambda b(s)) \tag{33}$$

$$+ \sum_{s=1}^{t_0} \epsilon(s)b(s) \prod_{\sigma=s+1}^t (1 - \lambda b(\sigma)) \tag{34}$$

$$+ \frac{\varepsilon\lambda}{3} \sum_{s=1}^t b(s) \prod_{\sigma=s+1}^t (1 - \lambda b(\sigma)). \tag{35}$$

Term (33) develops as

$$z(1) \prod_{s=1}^t (1 - \lambda b(s)) = z(1) e^{\sum_{s=1}^t \log(1 - \lambda b(s))} \leq z(1) e^{-\lambda \sum_{s=1}^t b(s)} \rightarrow 0,$$

by assumption (22c). Term (34) consists of t_0 terms (a finite number), each of which converges to 0 by the same argument as above. Hence, there exists t_1 such that for all $t \geq t_1$, we have

$$z(1) \prod_{s=1}^t (1 - \lambda b(s)) + \sum_{s=1}^{t_0} \epsilon(s)b(s) \prod_{\sigma=s+1}^t (1 - \lambda b(\sigma)) \leq \frac{2\varepsilon}{3}.$$

It can be shown that term (35) can be written as:

$$\frac{\varepsilon\lambda}{3} \sum_{s=1}^t b(s) \prod_{\sigma=s+1}^t (1 - \lambda b(\sigma)) = \frac{\varepsilon\lambda}{3} \frac{1}{\lambda} \left(1 - \prod_{s=1}^t (1 - \lambda b(s)) \right) \leq \frac{\varepsilon}{3}.$$

We have shown that for all $\varepsilon > 0$, there exists some $t_2 = t_0 \vee t_1$ such that for all $t > t_2$, we have $z(t+1) \leq \varepsilon$.

We thus obtain that when decomposing vector $\Psi(t)$ according to the eigenbasis of matrix Δ , one finds vanishing coordinates except along eigenvector e . Since the scalar product $e'\Psi(t)$ is always equal to $\sum_u g_u(t)$, as follows from (32), the announced result follows. \square

To conclude the proof of the theorem, note now that, by the previous lemma, and our specific choice of gain parameters $Y_u(t)$ in (17), for large enough t , Equation (16) reads in vector form

$$X(t+1) - X(t) = \frac{a(t)}{\max(1, \sum_{k,v} X_{vk}^2(t) + o(1))} [AX(t) - X(t)X'(t)AX(t) + o(\|X(t)\|) + \xi(t+1)]. \quad (36)$$

As is readily seen, this coincides with the update rule (23), except for the $o(\cdot)$ terms. The analysis of [2] for establishing convergence of (23) also applies in fact to its perturbed version (36), and Theorem 3 follows. \square